Summary of initial SSD review deliberations

Issues regarding multiple re-installation plan

SSD group – February 1, 2008

The committee assumes that the re-installation of the SSD in the summer of 2008 has been cancelled, and therefore the detector will be installed the earliest in 2009.

That statement is true.

The committee proposes to either install all 20 upgraded ladders in 2009 or wait with the full installation until 2010. The committee questions the motivation to run in 2009 with a hybrid system of 3 new and 17 old ladder configurations.

The current proposal is to install 3 ladders only, equipped with the upgraded electronics.

We believe the main physics argument is still the role the SSD will play in an integrated tracker after the HFT and IST have been installed. The only other physics driver could be usage of the SSD for displaced vertices in Run 10, which will presumably be the first low-energy heavy ion run.. Do the simulations in Figs. 1 and 3 take into account the primary vertex in the 'TPC only' curves? Although the physics performance for strange particles might be slightly improved the group needs to weigh the effort to have a working hybrid system of old and new electronics in Run 10 with the possibility of having a complete new system running in Run 11.

If the group decides to wait until the complete new systems can be installed, then resources that would be required to resolve the L2 abort problem could be re-directed to other issues.

The L2 problem may not go away as everyone seems to assume. There has been no definite decision to do away with L2 once the upgraded DAQ is installed.

We note that the present five year run straw-man proposes low energy heavy ion runs in Run 10 and Run 12, and a long pp run in Run 11. If the full SSD upgrade can not be completed for Run 10, then the detector presumably will not be fully utilized until Run 12.

[It is worthwhile to mention that the FGT also can utilize SSD+IST as helping in the pp W asymmetry program in specific eta region. Since FGT certainly is available for run 11 (possible 10) it would underscore the need for having the SSD as early as possible. Flemming V]

Detailed technical questions

1.) Please quantify the increase in power consumption due to the new electronics. Presently the detector consumes 400 W total power, with the ADC board being the hottest spot. The committee believes that this is close to the maximum power that can be handled by an air cooling system. By how much will the power increase and has the new air cooling system been designed accordingly?

Need Christophe to answer the power consumption question. The cooling system, as we know, has not yet been designed. A measurement of the power dissipated is part of the preliminary tests.

2.) The committee needs technical details about the new ADC board and the ADC's themselves (power consumption, readout speed, cost, test results etc.). Is there an ALICE web-page where all this information is collected and can be reviewed?

The ADC board has not yet been designed. As this board is different than the ALICE design, there is no documentation ready. We plan on using an ADC with a minimum speed of 5 MHz.

3.) The committee would like to see a detailed diagram of the RDO chain from wafer to ADC board to connection board to RDO board to DAQ.

This information is not yet available.

The link below represents the latest design.
 HYPERLINK "http://www-subatech.in2p3.fr/~electro/projets/star/star_readout/src/c_ssdU_d.jpg" http://www-subatech.in2p3.fr/~electro/projets/star/star_readout/src/c_ssdU_d.jpg

4.) How many ADC boards, how many connector boards, how many RDO boards, and how many VME-like RDO crates are necessary and where are they located? Specifically, is the readout fanned out to both sides of the TPC?

The readout is done from both ends of a ladder. One end is for the P face of the silicon wafer and the other is for the N face. They are independent of one another except for the common control electronics (bias, etc.).

Yes. One ADC board and one connector board at each end of each ladder
One VME-like RDO crate on each side.

Does each ladder have two ADC and two connector boards (one on each side) ?

Yes. One on each end.

In one part of the proposal you talk about a total of 4 RDO boards in another it says 8 RDO boards. Will there be RDO boards on both sides of the TPC (2 and 2 or 4 and 4) ?

4 on each side. [The reference to 4 RDOs is an editorial mistake – refers to a previous version of the system with less than adequate performance.]

Could the detector be read out to one side only? No.

Do the VME like crates that hold the RDO boards have to be mounted to the TPC wheel (old SVT location) or could the digital signal be routed to behind the pole-tip for better access and less radiation length in forward direction? (e.g. the new SVT RDO location would make more sense than the old location).

I believe the intended location for the RDO crates will be on the pole tip (old SVT location). [See pages 24 and 26 of the proposal.]

Contradiction with p 26 where the RDOs are located on the wheel.

The maximum cable run from the end of the ladder to the RDO will be determined in the preliminary tests.

Howard Matis 1/29/08 3:17 PM
Comment: Committee recommends to go to the pole-tip is possible.

Do these crates need cooling? If so, air or water?

I think so. Water cooling might be better, The idea is to utilise the cooling for the SVT boxes. Yes – depending on the power supply voltage, there will be 20 to 30W dissipated per RDO.

Is there a radiation length budget for each part of the RDO chain from wafer to DAQ ?

The radiation length budget doesn't change dramatically at the level of the barrel (wafer and electronics boards at the end of the ladder. You can find a reference to the radiation length of the current SSD at

http://www.star.bnl.gov/public/ssd/STAR_informatique/Geom/material.html

Are there discussions with the EEMC and the FGT groups on forward radiation length budget constraints?

5.) We propose to put buffers on the RDO boards independent of zero-suppression. In other words we propose to utilize dead-time managing scheme #1. Apparently this might force you to use 8 RDO boards rather than four boards, but we should not rely on low occupancy zero-suppression alone for dead time management.

Zero suppression will be done on the RDOs by comparing the pedestal-subtracted ADC count for each strip with a threshold value and sending the strip address and ADC count for each strip over threshold.
Choice of eight boards arises because 4 boards do not deliver sufficiently low dead time near 1kHz. The occupation of the DDL fibers with data is the reason for this

Does each RDO board need its own DDL link or could eight boards be served by four links via DDL jumpers?

We are not using a DDL jumper. We are not familiar with that term.

The committee expressed concern about the availability of PCIX slots for DDL link processing and therefore the smallest number of DDL links should be used.

Each DRORC (PCI-resident DDL receiver card) hosts 2 DDL links. A total of 4 such cards are required to service 8 RDOs. At worst this will require 2 PCs (cost $3K each). PCI-X is the most up-to-date standard in the PCI family of busses.

6.) We need to know the raw data volume without zero suppression and realistic anticipated occupancies in worst case scenarios (i.e. noisy central AuAu event) in order to assess the scope of the RDO chain.

Full detector readout without zero-suppression: 657 kB. An Au-Au central event should have 1% occupancy. With overhead for address encoding, this corresponds to 3% burden on the DDL link.

It should be clarified that the main reason to retain zero suppression is to relieve the PC of the memory-intensive chore of performing this in software.

It would be dangerous to assume that the effective occupancy will remain at 1 or 2%, since silicon wafers starting to misbehave can rapidly spoil this. This is the reason to retain multiple buffering in addition to the zero suppression on the RDO. We have always intended to use the multiple buffer technique.

7.) We propose to set the level of 'acceptable dead time' to less than 10% at 1 kHz trigger rate, because we should assume that all systems in DAQ1000 should be laid out for a 1 kHz trigger rate (even if this might not be fully achievable for the TPC). Figs.12/13 show that this acceptable dead time can be accomplished with the proposed readout chain. In order to evaluate the simulation, though, we would like the see at which trigger rate the linear increase in dead time breaks down.

The calculations have been carried out for trigger rates up to 5000 Hz.
The non-linear behavior of the (non zero suppressed) system becomes evident.
It is clear for the non-zero-suppressed system that the multi-buffered system's behavior approached that of the single buffer at a trigger rate of about 6000 Hz.

For the zero-suppressed system, the DDL transfer time is assumed to be 10 us (3% occupancy); the characteristic dead time corresponds to the FEE data transfer to the RDO (85 us), so the limiting behavior will be above 10 kHz.

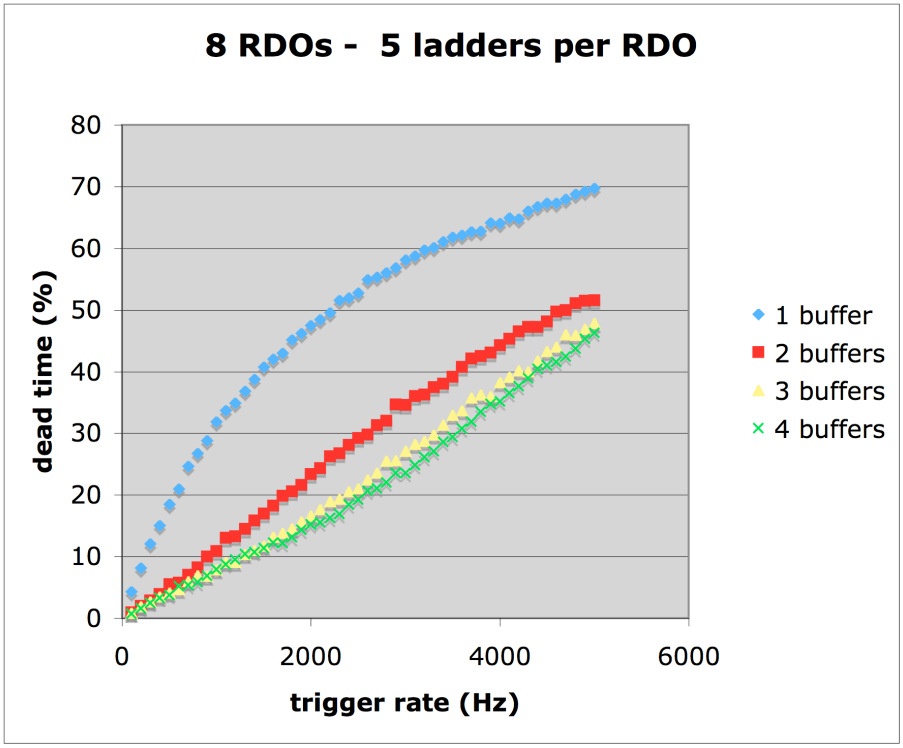
8 RDOs - 5 ladders per RDO
dead time (%)
80
70
60
50
40
30
20
10
0
0
2000
4000
6000
trigger rate (Hz)
1 buffer
2 buffers
3 buffers
4 buffers

8.) is there an agreement between ALICE and STAR regarding the usage of ALICE electronics design? How much of the design is actually being applied in the new SSD RDO chain? Please be detailed.

This impression should be corrected by stating that the front-end electronics to be used for the upgraded SSD readout has no relation to ALICE, and is based on commercial 8-channel ADC chips. The only connection to ALICE in the upgrade is the use of the DDL link (SIU, DRORC), which is already being used by the TPC for DAQ1000. The SIU, DIU, D-RORC are sold commercially to research organizations by an institute which is part of ALICE (Budapest). No additional management agreement is necessary.

9.) The cooling system integration between SSD and HFT needs to be better defined.

Yes. That is said explicitly in the proposal

Not only the cooling system integration, everything should be defined, the mechanical structure (the SSD mechanical structure remains the same) and the place where all the cables (power supply and readout) are routed.

10.) Regarding the upgrade of the TCD in STAR we recommend that the SSD group designs its system assuming the current TCD architecture. Any major TCD upgrade which will include hardware and protocol changes needs to be implemented in all STAR system, presumably through an additional mezzanine card.

We will be compatible with the STAR Trigger – whatever it is.

11.) We are concerned about shared real estate along the RDO chain with respect to HFT, IST, EEMC, FGT and TPC. A follow-up technical review should have a detailed look at cooling, cable routing and mounting structures.

We agree. A serious integration management effort is required. See answer 9).

The integration between HFT, SSD and FGT are being addressed by the HFT project, which has build in links explicitly to SSD, and implicitly to FGT (via eng sharing; I agree this is an important issue, and should be governed by change control. The layout of envelopes is in progress, but much work has to be done. The change control pathway should be agreed upon, particular now where the FGT in particular is moving forward. It is assumed that FGT MUST meet the requirement of readout + cooling via FGT (new west cone.

12.) The detector did not perform well in Run-7 (around 50% live). Many of the problems have likely occurred because of inhibited cooling. We agree that the cooling has to be made more robust, but radiation length constraints have to be met.

The inadequacy of the cooling was due to crimped hoses. When designing the cooling system, we will minimize the radiation length in the high rapidity region.

13.) Were ladders sent back to Nantes, repaired, and sent back to BNL ?

The ladders will be repaired in winter/spring 2008. Because the installation is cancelled in summer 2008, isn't it better to keep them at Subatech until the upgrade?

Was there any progress on the L2 problem?

Ladders are in Nantes. Diagnostics and repair are about to start (February, 2008). No progress yet on L2 problem. This will be tackled during the shutdown.

We will have a new DAQ when we install the upgraded section. The software and firmware will be completely different. We have learned that it is essesntial to test the hardware at BNL.

Cost and Schedule:

1.) How many FTE's can we expect from the four Nantes engineers and the one Nantes technician ?

The manpower available at Nantes is the following:

- Stephane Bouvier: project engineer : 10% (ladder repair, supervising the upgrade and testing the ladders)
- Gerard Guilloux: mechanical engineer : 5% (mounting/dismounting SSD and ladders)
- Christophe Renard : electronic engineer : less than 50% (upgrade designing)
- Louis-Marie Rigalleau: electronic technician: 30% (ladder repair and testing)

Is there an MoU in place that guarantees the labor resources in case Nantes receives the funds for its hardware components (~200 K)? We propose to have an MoU in place before reviewing the cost.

We plan to have a MoU once STAR agrees to fund this proposal. A MoU will be required before we start the project.

2.) We think the resource loading at Nantes and BNL is reasonable, and we recognize that the DoE needs to be made aware that the maintenance and operation of each new system will require an increase in STAR operations (e.g. 0.1 FTE for the SSD).

As part of the proposal, the US will take responsibility for the board layout and routing and PCB fabrication. The reason to take this action is to minimize the problems of funds transfer, and to transfer the expertise and responsibility for maintaining the electronics to the US.

3.) The technical roles of Nantes and BNL are well defined, but we do not understand the role of the LBL group on the proposal. We need to review a management structure proposal that shows the separation of work among the names on the proposal and shows a project leadership structure. It should also be shown how this structure interfaces with the HFT/IST project structure. Clearly many of the integration issues regarding mounting of the detector, the cooling system, routing of cables and cooling and placement of the RDO electronics will require detailed integration work and drawings. It would be useful to provide at least some schematic drawings to our committee, and again detail the interface and integration administration structure. A management structure should also specify the chain of command for all the engineers working on the SSD project alone.

4.) In order to assess the cost estimate in detail we need more information about the cost of the ALICE components. Generally the cost looks realistic, the resource loading is reasonable, and we believe the project can be completed for around 500 K as proposed.

There is as much detail in the proposal (spread sheets) as it is possible to generate.

5.) The schedule looks realistic as well, pending a the proposed funding profile. If the funding could be accelerated we don't see a bottleneck in the production in order to complete twenty instead of three ladders in 2009.

Full production cannot be started until a pilot run is installed and evaluated (3 ladders in 2009).